
AWS
re:Invent

DEM27-S

Drive transformation through machine learning with Amazon SageMaker

David Frigeri
Practice Area Lead
Slalom

Today’s key takeaways

- ✓ What is transformation?
- ✓ What value does Amazon SageMaker provide?
- ✓ What critical considerations help drive successful transformations?

aws | slalom

About us

An overview of Slalom and AWS

September 2019

Why AWS & Slalom?

Because our cultures, values, and delivery approaches are uniquely well-aligned

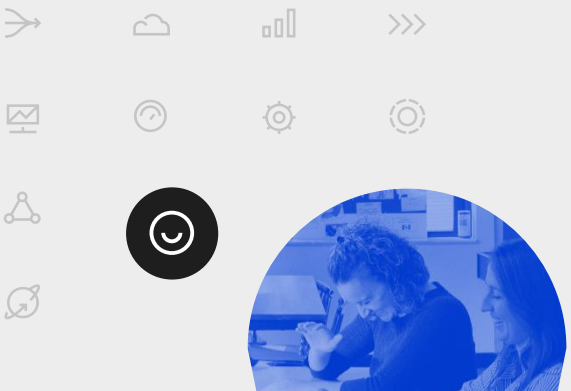

Slalom

Customer love

Do what is right, always

Drive connection + teamwork

Take ownership. Get it done.

Inspire passion + adventure

Focus on outcomes

Celebrate authenticity

Fuel growth + innovation

Stay humble + curious

Build + shape a better future

Smile!

AWS

Customer obsession

Ownership

Invent and simplify

Are right, a lot

Learn and be curious

Hire and develop the best

Insist on the highest standards

Think big

Bias for action

Frugality

Earn trust

Have backbone; disagree and commit

Deliver results

aws partner network

Get transformative, ongoing value from the AWS Cloud

We've established ourselves as leaders in cloud transformation by helping our clients maximize the value of the AWS Cloud, from planning to migration and all that comes next.

Premier
Consulting Partner

1250+
AWS certifications

670+
certified AWS consultants

618
projects delivered through Oct. 2019

AWS competencies

- Data & analytics
- DevOps
- Education
- Financial services
- Government
- Machine learning
- Microsoft workloads
- Migration
- Mobile
- Public sector partner

aws | slalom

Why cloud and machine learning?

September 2019

Machine learning (ML) in the cloud

Model modernization

Optimize model development, including third-generation modern languages and APIs

Create a competitive advantage, including model-driven competitors

Outsource commodity infrastructure and platform services

Reduce downtime and increase productivity

Support international market expansion

Increase responsiveness, flexibility, and scalability

Organizational evolution

Remediate legacy code debt, and provide a catalyst for change

Unlock headcount to fuel transformative model development

Increase speed to market for changes, features, and products

Reduce IT costs, transitioning from capex to opex

ML drives value early and often

53%

Companies who don't have a clear understanding of how AI or ML could benefit their businesses

~TechRepublic

45%

Estimated percent of work activities that can be automated with current ML capabilities

~US Bureau of Labor Statistics

60%

Percent of occupations that have at least 30% of activities that could be automated.

~US Bureau of Labor Statistics

8.4%

Increase in number of Fortune 1000 companies reporting $500M+ in investments from 2018 (12.7%) to 2019 (21.1%)

~MIT Sloan Management Review

91%

Fortune 1000 companies citing the need to transform their businesses to become increasingly nimble

~MIT Sloan Management Review

aws | slalom

Cloud + machine learning

We need managed ML

Businesses are becoming more sophisticated consumers of ML, and expectations are rising rapidly

Better alignment

Connect ML initiatives to corporate strategy

Deliver the right insights

- To the right people
- At the right time
- In the right form

Speed to answers

Knowing before competitors and customers is a strategic capability

Build the flywheel where every project is completed faster than before

Use the broader world of big data

See insights from beyond what's isolated inside four walls of a business, which means more data sources and larger datasets

Expand the accessibility insights

Predict AND impact outcomes

Increase interactions

Create simulation capabilities

Drive continuous learning

Understand actions likely to result in desired outcome

Cloud + machine learning

Amazon SageMaker enables business transformation

Amazon SageMaker is an end-to-end modular service used to build, train, and deploy ML models quickly and at scale

- Amazon SageMaker provides common ML algorithms optimized to run efficiently against extremely large data in a distributed environment
- Models deploy into secure and scalable environments, launching with a single click from the Amazon SageMaker console
- Data scientists and developers can quickly and easily build, train, and then deploy ML models to a production-ready hosted environment
- Models get to production faster with much less effort, with lower cost (which enables transformation), and with less risk

aws | slalom

aws | slalom

What to migrate?

Best fit models or use cases for a fully managed ML solution

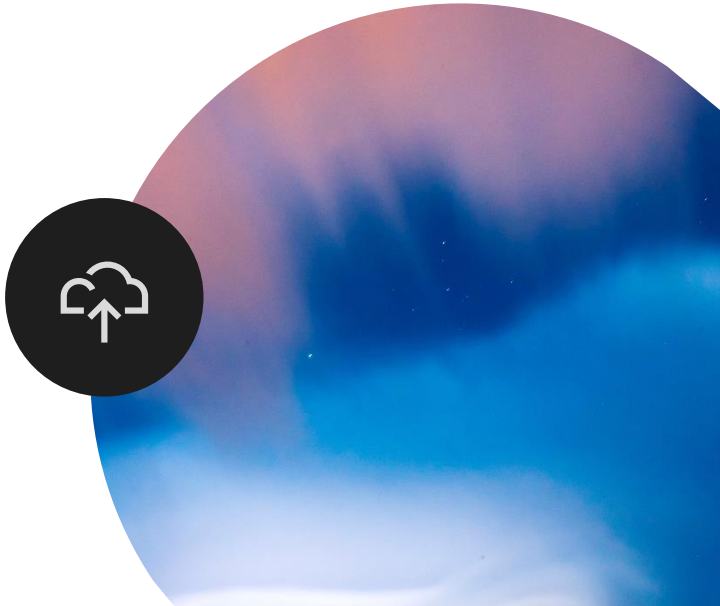

September 2019

What to migrate?

Define your “why” with a few questions

1

What are the key business outcomes you aim to achieve?

Establish the top line and strategic reasons for migration to Amazon SageMaker

2

What would happen if you didn’t migrate to Amazon SageMaker?

Define the burning platform; customer, speed, and cost are common reasons

3

Can you quantify the value of your migration to Amazon SageMaker?

Develop a strong business case for your benefits, including key business imperatives and prioritized backlog use cases

4

How will you define the success of your efforts?

Establish targeted, measurable outcomes for balancing accuracy and continuity, business integration, utilization, and self and service adoption

What to migrate?

Drivers behind what models to migrate

Customers take the following things into consideration when determining what models to migrate

Old models

Models that are old, written in proprietary code, and maybe the original data scientist is no longer available to maintain

To-be-developed

Models that are on the roadmap and yet to be developed:
"We have been wanting to build this"

Manual processes

Models that need to be manually processed on a daily and weekly basis

Storage limits

Models that require significantly more data and for which cloud storage is an imperative

External factors

Need to reduce or reallocate costs, models currently running in data centers due to be shuttered, corporate acquisitions or other reorganizations, employee attrition, models that are owned by data scientists who have been recently trained or hired with modern cloud and development capabilities (e.g., Python)

What to migrate?

Options for migrating analytics workloads to AWS

Customers tend to have multiple interrelated reasons for migrating, including cost reduction, modernizing, or skill set availability

- Forklift entire environment onto AWS; e.g., Cloudera (Hadoop and Spark implementations)
- Move analytics server onto Amazon EC2 and pull data from Amazon S3; e.g., Teradata Vantage
- Package legacy models in ONNX, pull data from Amazon S3, and execute in Amazon SageMaker; e.g., SAS
- Rebuild and/or update models in Python, pull data from Amazon S3, and execute in Amazon SageMaker

How to decide

Stages	Description
Pre-work	• Align with business on vision and communication plan • Agree on success criteria; for example, is maintaining continuity more important than trying to improve accuracy? • Educate the business with examples of how other organizations are taking advantage of AWS machine learning • Identify and prioritize those in the business willing to partner on and support the migration

Model assessment and identification	• Evaluate current performance of models—how well are they documented? Do you see them running 12 months from now?
	  • Should there be a more sophisticated algorithm employed?
	  • Unclear what the original intention of the model was
	  • Unable to identify the original data
	• Are there models out of current coding standards, such as
	  • Typo-causing issues
	  • Hard code date and others
	  • Unsupported code; e.g., Python 2.x
	  • Not properly treating NAs
	• Do you have multiple models basically predicting the same thing, and there's opportunity for consolidating and simplifying?
	• Are there models outputting .CSV but should be accessed via API or data visualization?

Data readiness	• Are any models not loading complete input data or loading extra data? • Should new data sources and features be added to the model? • Is the model's data source scheduled to be moved to Amazon S3?

aws | slalom

Customer story

Let’s talk about Avis

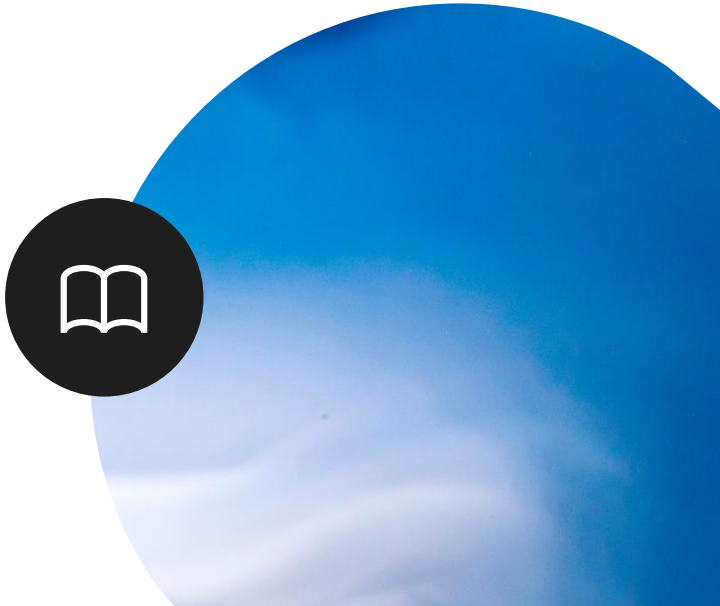

September 2019

Customer story

Our why

The Avis Budget Group, Slalom, and AWS teams had several use cases to build on the connected car platform

Use case #1

Continually rent out (re-purchase/lease) program cars

Use case #2

Load balance mileage across risk cars

aws | slalom

Customer story

Mileage optimization ML architecture

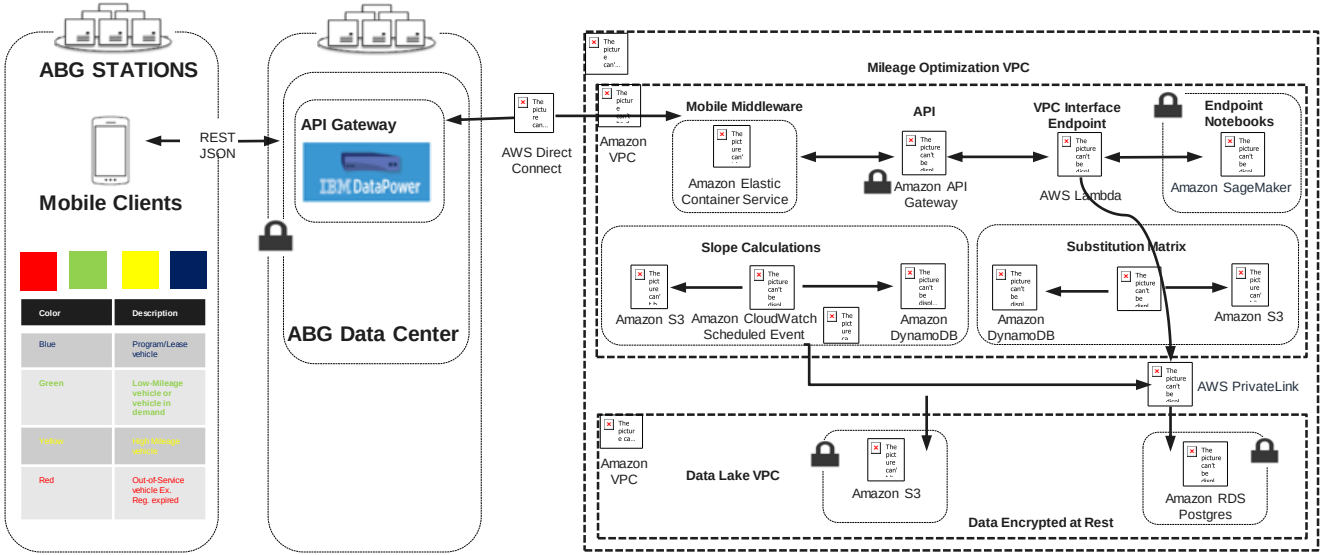

aws | slalom

Key points

Leverages the AWS connected vehicle framework as a foundation and expands on it for key use cases

Takes the modern architecture stack into consideration by including both on-premises and AWS architecture

Designed for real-time analytics and for operational efficiency

Customer story

Our model development followed an agile, sprint-based process

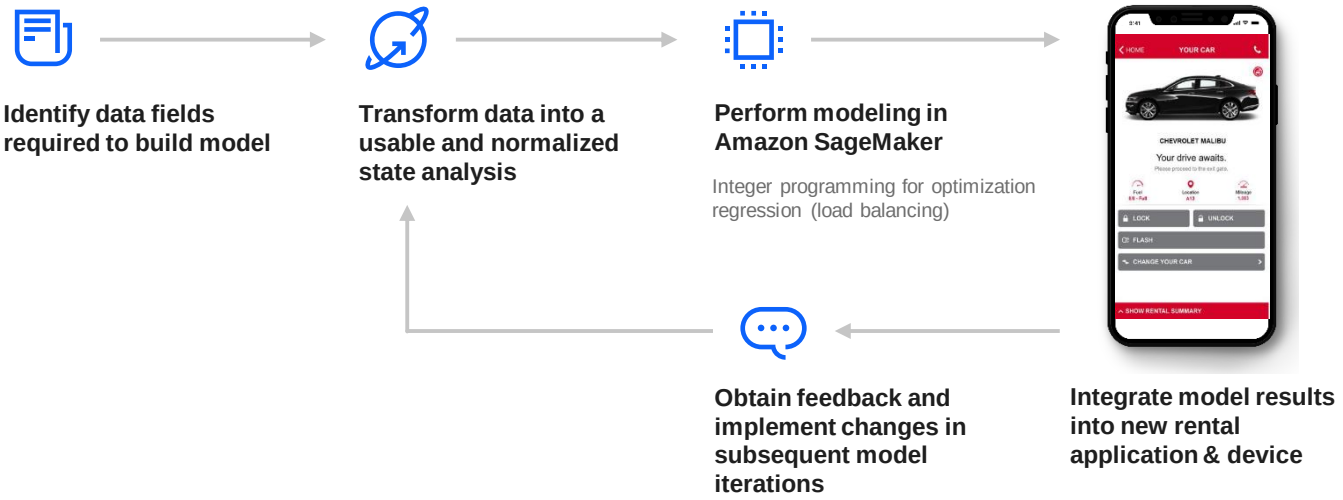

Customer story

Mileage optimization ML architecture

Component	Usage
AWS Lambda	Invokes Amazon SageMaker endpoint and sends input data Receives optimized result from Amazon SageMaker Stores response in Amazon S3 bucket (for future model analysis) Sends response back to calling API
Amazon SageMaker	Imports relevant Python libraries Generates supply and demand for the specified time horizon Incorporates upgrade and substitution logic Simulate slope and average miles per day (MPD) Calculates “ideal” mileage for any particular vehicle Calculates ideal vs. actual ratio

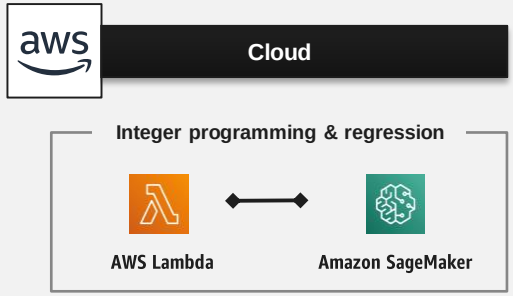

Amazon SageMaker algorithms

Integer programming
Mathematical optimization or feasibility program in which some or all of the variables are restricted to be integers; used to assign optimal mileage to vehicles

Regression analytics
Predicting load balancing across vehicles

Customer story

The model takes key operational components into consideration for deployment

Feature	Description
Cars on readyline	Rentable cars that are already parked in the readyline
Demand horizon	Helps determine if there is a demand (reservation) for a particular car class in the next horizon (particular number of hours) of time
Primary vs. exclusive	The cars are divided into two cohorts, where substitutions can happen within a cohort
Exclusive cars by class	Demand for an exclusive class car can only be satisfied by that car class
Substitution	Demand for a primary class car can be satisfied by other primary class cars
Contention	Satisfying a reservation with an upgrade car class if a better car is available

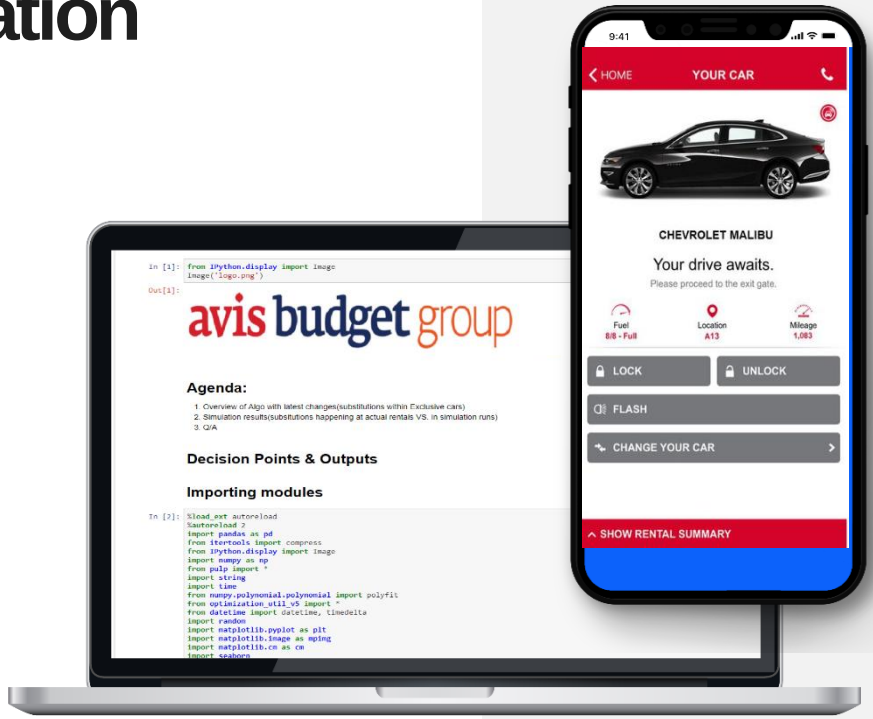

Please complete the session survey in the mobile app.